

IN THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

- 1 1. (Currently Amended) A mass storage controller system, comprising:
2 a plurality of controllers for controlling an array of storage devices, each of the
3 plurality of controllers comprising:
4 a CPU for controlling the operation of a controller;
5 program memory, coupled to the CPU, for storing program instructions and
6 variables for the operation of the CPU; and
7 cache memory, coupled to the CPU, for storing information related to the
8 array of storage devices;
9 wherein a controller of the plurality of controllers initiates a task to be performed, the
10 controller initiating the task establishes a task coordination data object shared by the plurality
11 of controllers, wherein the task coordination data object represents discrete partitions of task
12 instructions of the task that may be performed independently to be performed and states for
13 each partition of task instructions to allow the plurality of controllers to cooperate in the
14 execution of the task, controllers performing steps of the task independently of other of the
15 plurality of controllers, and wherein a free controller of the plurality of controllers selects a
16 partition of task instructions of the task available for processing as indicated by the states for
17 each partition of task instructions.
- 1 2. (Original) The mass storage controller system of claim 1, wherein the
2 state indicates whether a partition is READY, IN PROGRESS, or COMPLETE.
- 1 3. (Original) The mass storage controller system of claim 2, wherein a
2 controller selects a partition by examining the partitions in a READY state and selecting at
3 least one partition in the READY state to operate on.
- 1 4. (Original) The mass storage controller system of claim 3, wherein a
2 partition is in an IN PROGRESS state during processing.

1 5. (Original) The mass storage controller system of claim 4, wherein a
2 controller sets the partition selected for processing to a COMPLETE state upon completion
3 of processing for a partition.

1 6. (Original) The mass storage controller system of claim 1, wherein a
2 controller selects a partition by examining the partitions in a READY state and selecting at
3 least one partition in the READY state to operate on.

1 7. (Original) The mass storage controller system of claim 1, wherein a
2 partition is in an IN PROGRESS state during processing.

1 8. (Original) The mass storage controller system of claim 1, wherein a
2 controller sets the partition selected for processing to a COMPLETE state upon completion
3 of processing for a partition.

1 9. (Original) The mass storage controller system of claim 1, wherein the
2 states provide a semaphore-mechanism for allowing a controller to ascertain whether to
3 acquire control over a partition.

1 10. (Original) The mass storage controller system of claim 1, wherein the
2 initiating controller is notified when all partition states are COMPLETE and performs
3 whatever completion actions are required.

1 11. (Original) The mass storage controller system of claim 1, wherein the task
2 coordination data object includes information about an operation to be performed and a data
3 set to be operated on.

1 12. (Currently Amended) A mass storage array subsystem, comprising:
2 a plurality of storage devices;
3 a backplane, coupled to the plurality of storage devices, adapted to couple to said
4 plurality of storage devices; and
5 a plurality of controllers, coupled to the backplane, for controlling the plurality of
6 storage devices, the plurality of controllers having a first interface coupled to a host system
7 and a second interface coupled to said backplane to communicate with said plurality of
8 storage devices;
9 wherein each of the plurality of controllers comprise a CPU for controlling the
10 operation of a controller, program memory for storing program instructions and variables for
11 the operation of the CPU and cache memory for storing information related to the array of
12 storage devices, and wherein a controller of the plurality of controllers initiates a task to be
13 performed, the controller initiating the task establishes a task coordination data object shared
14 by the plurality of controllers, wherein the task coordination data object represents discrete
15 partitions of task instructions of the task that may be performed independently to be
16 performed and states for each partition of task instructions to allow the plurality of
17 controllers to cooperate in the execution of the task, controllers performing steps of the task
18 independently of other of the plurality of controllers, and wherein a free controller of the
19 plurality of controllers selects a partition of task instructions of the task available for
20 processing as indicated by the states for each partition of task instructions.

1 13. (Original) The mass storage array subsystem of claim 12, wherein the
2 state indicates whether a partition is READY, IN PROGRESS, or COMPLETE.

1 14. (Original) The mass storage array subsystem of claim 13, wherein a
2 controller selects a partition by examining the partitions in a READY state and selecting at
3 least one partition in the READY state to operate on.

1 15. (Original) The mass storage array subsystem of claim 14, wherein a
2 partition is in an IN PROGRESS state during processing.

1 16. (Original) The mass storage array subsystem of claim 15, wherein a
2 controller sets the partition selected for processing to a COMPLETE state upon completion
3 of processing for a partition.

1 17. (Original) The mass storage array subsystem of claim 12, wherein a
2 controller selects a partition by examining the partitions in a READY state and selecting at
3 least one partition in the READY state to operate on.

1 18. (Original) The mass storage array subsystem of claim 12, wherein a
2 partition is in an IN PROGRESS state during processing.

1 19. (Original) The mass storage array subsystem of claim 12, wherein a
2 controller sets the partition selected for processing to a COMPLETE state upon completion
3 of processing for a partition.

1 20. (Original) The mass storage array subsystem of claim 12, wherein the
2 states provide a semaphore-mechanism for allowing a controller to ascertain whether to
3 acquire control over a partition.

1 21. (Original) The mass storage array subsystem of claim 12, wherein the
2 initiating controller is notified when all partition states are COMPLETE and performs
3 whatever completion actions are required.

1 22. (Original) The mass storage array subsystem of claim 12, wherein the
2 task coordination data object includes information about an operation to be performed and a
3 data set to be operated on.

1 23. (Currently Amended) A method for cooperative distributed task management
2 in a storage subsystem with multiple controllers, comprising:
3 initiating by an initiating controller a task to be performed;
4 establishing by the initiating controller a task coordination data object shared by the
5 multiple controllers, wherein the task coordination data object represents discrete partitions
6 of task instructions of the task to be performed that may be performed independently and
7 states for each partition of task instructions to allow the plurality of controllers to cooperate
8 in the execution of the task, the plurality of controllers performing steps of the task
9 independently of other of the plurality of controllers; and
10 selecting by a free controller a partition of task instructions of a task available for
11 processing as indicated by the states for each partition of task instructions.

1 24. (Original) The method of claim 23 further comprising indicating a state of
2 a partition as being READY, IN PROGRESS, or COMPLETE.

1 25. (Original) The method of claim 24 wherein the selecting by a free
2 controller is performed by examining the partitions in a READY state and selecting at least
3 one partition in the READY state to operate on.

1 26. (Original) The method of claim 25, wherein a partition is in an IN
2 PROGRESS state during processing.

1 27. (Original) The method of claim 26 further comprising setting by a
2 controller a partition selected for processing to a COMPLETE state upon completion of
3 processing for the partition.

1 28. (Original) The method of claim 23, wherein the selecting by a free
2 controller is performed by examining the partitions in a READY state and selecting at least
3 one partition in the READY state to operate on.

1 29. (Original) The method of claim 23, wherein a partition is in an IN
2 PROGRESS state during processing.

1 30. (Original) The method of claim 23 further comprising setting by a
2 controller a partition selected for processing to a COMPLETE state upon completion of
3 processing for the partition.

1 31. (Original) The method of claim 23 wherein the states provide a
2 semaphore-mechanism for allowing a controller to ascertain whether to acquire control over
3 a partition.

1 32. (Original) The method of claim 23 further comprising notifying the
2 initiating controller when all partition states are complete and performing completion actions
3 that are required.

1 33. (Original) The method of claim 23, wherein the task coordination data
2 object includes information about an operation to be performed and a data set to be operated
3 on.

1 34. (Currently Amended) An article of manufacture comprising a program
2 storage medium readable by a computer, the medium tangibly embodying one or more
3 programs of instructions executable by the computer to perform a method for cooperative
4 distributed task management in a storage subsystem with multiple controllers, the method
5 comprising:
6 initiating by an initiating controller a task to be performed;
7 establishing by the initiating controller a task coordination data object shared by the
8 multiple controllers, wherein the task coordination data object represents discrete partitions
9 of task instructions of the task to be performed that may be performed independently and
10 states for each partition of task instructions to allow the plurality of controllers to cooperate
11 in the execution of the task, the plurality of controllers performing steps of the task
12 independently of other of the plurality of controllers; and
13 selecting by a free controller a partition of task instructions of a task available for
14 processing as indicated by the states for each partition of task instructions.

1 35. (Original) The article of manufacture of claim 34 further comprising
2 indicating a state of a partition as being READY, IN PROGRESS, or COMPLETE.

1 36. (Original) The article of manufacture of claim 35 wherein the selecting by
2 a free controller is performed by examining the partitions in a READY state and selecting at
3 least one partition in the READY state to operate on.

1 37. (Original) The article of manufacture of claim 36, wherein a partition is in
2 an IN PROGRESS state during processing.

1 38. (Original) The article of manufacture of claim 37 further comprising
2 setting by a controller a partition selected for processing to a COMPLETE state upon
3 completion of processing for the partition.

1 39. (Original) The article of manufacture of claim 34, wherein a free
2 controller selects a partition by examining the partitions in a READY state and selecting at
3 least one to operate on.

1 40. (Original) The article of manufacture of claim 34 further comprising
2 notifying the initiating controller is notified when all partition states are complete and
3 performs whatever completion actions required.

1 41. (Original) The article of manufacture of claim 34, wherein the task
2 coordination data object includes information about an operation to be performed and a data
3 set to be operated on.

1 42. (Currently Amended) A mass storage controller system, comprising:
2 a plurality of intermediate controller means for controlling an array of storage
3 devices, each of the plurality of controllers comprising:
4 CPU means for controlling the operation of a controller;
5 program memory means, coupled to the CPU means, for storing program
6 instructions and variables for the operation of the CPU; and
7 cache memory means, coupled to the CPU means, for storing information
8 related to the array of storage devices;
9 wherein an intermediate controller means of the plurality of intermediate controller
10 means initiates a task to be performed, the intermediate controller means initiating the task
11 establishes a task coordination data object shared by the plurality of intermediate controller
12 means controller means, wherein the task coordination data object represents discrete
13 partitions of task instructions that may be performed independently of the task to be
14 performed and states for each partition of task instructions to allow the plurality of
15 controllers to cooperate in the execution of the task, the plurality of controllers performing
16 steps of the task independently of other of the plurality of controllers, and wherein a free
17 intermediate controller means selects a partition of task instructions of the task available for
18 processing as indicated by the states for each partition of task instructions.